Independent Update Reflections on Interdependent Database Views

Stephen J. Hegner

Umeå University, Department of Computing Science
SE-901 87 Umeå, Sweden
hegner@cs.umu.se http://www.cs.umu.se/~hegner

Abstract. The problem of identifying suitable view-update strategies is typically addressed in the context of a single view. However, it is often the case that several views must co-exist; the challenge is then to find strategies which allow one view to be updated without affecting the other. The classical constant-complement strategy can provide a solution to this problem; however, both the context and the admissible updates are quite limited. In this work, the updates which are possible within this classical approach are extended substantially via a technique which considers only the states which are reachable from a given initial configuration. The results furthermore do not depend upon complementation, and thus are readily extensible to settings involving more than two views.

1 Introduction

Both views and updates are fundamental to a comprehensive database system. Consequently, the problem of how to support updates to views has been studied extensively. Most work addresses this problem in the context of a single view, including the classical approach via the relational algebra [10, 18, 19, 7, 8], the more recent approach based upon *database repairs* [1, 3, 2], and work which bridges these two approaches [12]. However, in some situations a number of distinct yet interdependent views of the same main schema must co-exist. Often, the access rights to these views differ, so that a user or access rôle [4, 21] which has access to one view may not even be allowed to read, much less update, another. In such a setting, it is important to identify those updates which are possible to a given view Γ without requiring any access to the other views, for reading or for writing. This may be recaptured succinctly in terms of two independence conditions. First of all, whether or not an update to Γ is to be allowed at all should be independent of the states of the other views. This is called *context independence*. Second, the reflection to the main schema of the update to the selected view must not require a change of the state of any of the other views. This is called *propagation independence* or *locality of effect*. In the presence of these two forms of independence, an update may be made to the given view Γ without knowledge about the states of the other views beyond that which is already known in Γ, and the result of the update to Γ will not be visible in any of the other views. Applications in which such independence is central, and which have motivated this work, include component-based architectures [24, 23, 13, 16], update by cooperation [17], and models of data objects for transactions [15].

For the case of two views, the classical constant-complement approach [6, 11] already provides a very elegant solution in the situations to which it applies. Unfortunately, it imposes conditions which are often too strong to be of use, as illustrated by the following examples.

Let $\mathbf{E}_0$ be the relational schema consisting of the single relation symbol $R[ABC]$, constrained by the functional dependency (FD) $B \to C$. Define $\Pi_{AB}^{\mathbf{E}_0} = (\mathbf{E}_0^{AB}, \pi_{AB}^{\mathbf{E}_0})$ to be the view whose schema $\mathbf{E}_0^{AB}$ contains the single relation symbol is $R_{AB}[AB]$ and whose morphism $\pi_{AB}^{\mathbf{E}_0}$ is the projection of $R[ABC]$ onto $R_{AB}[AB]$. Define $\Pi_{BC}^{\mathbf{E}_0} = (\mathbf{E}_0^{BC}, \pi_{BC}^{\mathbf{E}_0})$ analogously. Let $\mathsf{LDB}(\mathbf{E}_0)$ denote the set of all *legal databases* of $\mathbf{E}_0$; that is, the set of all relations on $R[ABC]$ which satisfy the FD $B \to C$. Define $\mathsf{LDB}(\mathbf{E}_0^{AB})$ and $\mathsf{LDB}(\mathbf{E}_0^{BC})$ similarly, as the legal databases of the corresponding view schemata. Define the the *decomposition mapping* $\pi_{AB}^{\mathbf{E}_0} \times \pi_{BC}^{\mathbf{E}_0} : \mathsf{LDB}(\mathbf{E}_0) \to \mathsf{LDB}(\mathbf{AB}_0^E) \times \mathsf{LDB}(\mathbf{BC}_0^E)$ on elements by $M \mapsto (\pi_{AB}^{\mathbf{E}_0}(M), \pi_{BC}^{\mathbf{E}_0}(M))$.

Let $u_1 = (N_1, N_1')$ be any update on $\Pi_{AB}^{\mathbf{E}_0}$, with N_1 representing the view state before the update operation and N_2 the state afterwards. A *reflection* of u_1 to $\mathbf{E}_0$ is any $(M_1, M_2) \in \mathsf{LDB}(\mathbf{D}) \times \mathsf{LDB}(\mathbf{D})$ with $\pi_{AB}^{\mathbf{E}_0}(M_1) = N_1$ and $\pi_{AB}^{\mathbf{E}_0}(M_2) = N_1'$. This update is *propagation independent* with respect to $\Pi_{BC}^{\mathbf{E}_0}$, or *keeps* $\Pi_{BC}^{\mathbf{E}_0}$ *constant*, if $\pi_{BC}^{\mathbf{E}_0}(M_1) = \pi_{BC}^{\mathbf{E}_0}(M_2)$.

In this example, the set of all updates on $\Pi_{AB}^{\mathbf{E}_0}$ which keep state of $\Pi_{BC}^{\mathbf{E}_0}$ constant has a very simple characterization; namely, it is precisely the set of all updates on R_{AB} which keep the projection onto B fixed. Similarly, the set of all updates on $\Pi_{BC}^{\mathbf{E}_0}$ with $\Pi_{AB}^{\mathbf{E}_0}$ constant is precisely the set of all updates on R_{BC} which keep the projection onto B fixed. The view $\Pi_{B}^{\mathbf{E}_0}$ of $\mathbf{E}_0$ which is the projection onto B, is called the *meet* of $\Pi_{AB}^{\mathbf{E}_0}$ and $\Pi_{BC}^{\mathbf{E}_0}$. For both $\Pi_{AB}^{\mathbf{E}_0}$ and $\Pi_{BC}^{\mathbf{E}_0}$, the updates which are propagation independent are precisely those which keep the meet view $\Pi_{B}^{\mathbf{E}_0}$ constant. Thus, whether or not an update to either view is possible without modifying the state of the other is a property of the state of that view alone, and does not require knowledge further knowledge of the state of $\mathbf{E}_0$; i.e., it exhibits context independence. For a more thorough presentation of these ideas in the context of update via constant complement, see [11, 1.2].

Pairs of views are not always so well behaved. Let $\mathbf{E}_1$ be identical to $\mathbf{E}_0$, save that it is governed by the additional FD $A \to C$, and let $\Pi_{AB}^{\mathbf{E}_1}$ and $\Pi_{BC}^{\mathbf{E}_1}$ be defined analogously to $\Pi_{AB}^{\mathbf{E}_0}$ and $\Pi_{BC}^{\mathbf{E}_0}$. The set of updates on $\Pi_{AB}^{\mathbf{E}_1}$ which are propagation independent with respect to $\Pi_{BC}^{\mathbf{E}_1}$ is not independent of the particular state of $\Pi_{BC}^{\mathbf{E}_1}$. For example, consider the two states $M_{10} = \{R(\mathrm{a}_1, \mathrm{b}_1, \mathrm{c}_1), R(\mathrm{a}_2, \mathrm{b}_2, \mathrm{c}_1)\}$ and $M_{10'} = \{R(\mathrm{a}_1, \mathrm{b}_1, \mathrm{c}_1), R(\mathrm{a}_2, \mathrm{b}_2, \mathrm{c}_2)\}$ in $\mathsf{LDB}(\mathbf{E}_1)$. Then $\pi_{AB}^{\mathbf{E}_1}(M_{10}) = \pi_{AB}^{\mathbf{E}_1}(M_{10'}) = \{R_{AB}(\mathrm{a}_1, \mathrm{b}_1), R_{AB}(\mathrm{a}_2, \mathrm{b}_2)\}$. The view update which replaces $\{R_{AB}(\mathrm{a}_1, \mathrm{b}_1), R_{AB}(\mathrm{a}_2, \mathrm{b}_2)\}$ with $\{R_{AB}(\mathrm{a}_1, \mathrm{b}_1), R_{AB}(\mathrm{a}_1, \mathrm{b}_2)\}$ on $\Pi_{AB}^{\mathbf{E}_1}$ has a reflection which keeps the state of $\Pi_{BC}^{\mathbf{E}_1}$ constant from M_{10} but not from $M_{10'}$. Thus, this view update does not exhibit context independence.

The key difference between $\mathbf{E}_0$ and $\mathbf{E}_1$ is that in the former the governing FDs embed into the views, while in the latter they do not. That this properly is necessary and sufficient to guarantee that the view updates which are possible while keeping a second view constant are independent of the state of the other view was first presented in [22, Thm. 2], and in a much more general context in [11, Prop. 2.17].

The conventional wisdom is that context-independent updates to views such as $\Pi_{AB}^{\mathbf{E}_1}$ are not possible, because checking the FD $A \to C$ requires access to both views. While this is true if one insists upon characterizing the allowable view updates as those which keep a meet view constant, it is nevertheless possible to support weaker, but still very useful, forms of context and propagation independence in such settings. It is the main goal of this paper to develop such notions of independence.

Given $N_1 \in \mathsf{LDB}(\mathbf{E}_1^{AB})$, let $\pi_B(N_1)$ denote $\{\mathrm{b} \mid (\exists t \in N_1)(t[B] = b\}$, that is, the set of all values for attribute B which occur in some tuple of N_1, and let $\equiv_{\langle B,A\rangle}^{N_1}$ denote the equivalence relation on $\pi_B(N_1)$ which identifies two B-values iff they share a common value for attribute A. Thus, $\mathrm{b}_1 \equiv_{\langle B,A\rangle}^{N_1} \mathrm{b}_2$ iff there are tuples $t_1, t_2 \in N_1$ with $t_1[B] = \mathrm{b}_1$, $t_2[B] = \mathrm{b}_2$, and $t_1[A] = t_2[A]$. It is not difficult to see that any view update (N_1, N_1') to $\Pi_{AB}^{\mathbf{E}_1}$ for which $\pi_B(N_1) = \pi_B(N_1')$ and for which $\equiv_{\langle B,A\rangle}^{N_1'} \subseteq \equiv_{\langle B,A\rangle}^{N_1}$ cannot lead to a violation of the FD $A \to C$ as long as the state of $\Pi_{BC}^{\mathbf{E}_1}$ is held constant in the reflection. For example, if the current state of $\Pi_{AB}^{\mathbf{E}_1}$ is $N_{11} = \{R_{AB}(\mathrm{a}_1, \mathrm{b}_1), R_{AB}(\mathrm{a}_2, \mathrm{b}_2), R_{AB}(\mathrm{a}_2, \mathrm{b}_3)\}$, then the update to the new state $N_{11'} = \{R_{AB}(\mathrm{a}_1, \mathrm{b}_1), R_{AB}(\mathrm{a}_2, \mathrm{b}_2), R_{AB}(\mathrm{a}_3, \mathrm{b}_3)\}$, as well as to the new state $N_{11''} = \{R_{AB}(\mathrm{a}_1, \mathrm{b}_1), R_{AB}(\mathrm{a}_3, \mathrm{b}_2), R_{AB}(\mathrm{a}_3, \mathrm{b}_3)\}$, cannot possibly result in a violation of $A \to C$, as long as the state of $\Pi_{BC}^{\mathbf{E}_1}$ is held constant, regardless of what that state is. In other words, limiting the view updates to those which satisfy these properties results in a strategy which is both context and propagation independent. A similar argument holds for updates on $\Pi_{BC}^{\mathbf{E}_1}$. For any $N_2 \in \mathsf{LDB}(\mathbf{E}_1^{BC})$, let $\equiv_{\langle B,C\rangle}^{N_2}$ denote the equivalence relation on $\pi_B(N_2)$ which identifies two B-values if they share a common value for attribute C. Now, any view update (N_2, N_2') with $\pi_B(N_2) = \pi_B(N_2')$ and for which $\equiv_{\langle B,C\rangle}^{N_2} \subseteq \equiv_{\langle B,C\rangle}^{N_2'}$ has a reflection with constant $\Pi_{AB}^{\mathbf{E}_1}$ which is both context and propagation independent. Furthermore, these updates may be made to $\Pi_{AB}^{\mathbf{E}_1}$ and $\Pi_{BC}^{\mathbf{E}_1}$ independently of each other without violating any integrity constraints. The compromise, relative to that of the views of $\mathbf{E}_0$, is that the allowable updates are with respect to a given initial context $(N_1, N_2) \in \mathsf{LDB}(\mathbf{E}_1^{AB}) \times \mathsf{LDB}(\mathbf{E}_1^{BC})$. In the case of $\mathbf{E}_0$, the identification of independent updates to $\Pi_{AB}^{\mathbf{E}_0}$ requires no knowledge of the state of $\Pi_{BC}^{\mathbf{E}_0}$. In the case of $\mathbf{E}_1$, knowledge that the state of each view is the result of context-independent updates from a consistent initial state is necessary. Furthermore, each view must know its image of that initial state. Thus, $\Pi_{AB}^{\mathbf{E}_1}$ must know N_1 and $\Pi_{BC}^{\mathbf{E}_1}$ must know N_2 (but $\Pi_{AB}^{\mathbf{E}_1}$ need not know N_2 and $\Pi_{BC}^{\mathbf{E}_1}$ need not know N_1).

There is a further improvement which may be made. Note that the set of allowable updates in this example is not symmetric. For example, updating the state of $\Pi^{\mathbf{E}_1}_{AB}$ from N_{11} to $N_{11'}$ is always admissible, but the reverse, from $N_{11'}$ to N_{11} is not, since the latter may lead to a violation of $A \rightarrow C$ for certain compatible states of $\Pi^{\mathbf{E}_1}_{BC}$. Nevertheless, for any $N_{12} \in \mathsf{LDB}(\mathbf{E}_1^{BC})$ which is compatible with N_{11} in the sense that they arise from a common $M \in \mathsf{LDB}(\mathbf{E}_1)$, this update is reversible. In fact, it remains reversible if the only updates to $\Pi^{\mathbf{E}_1}_{BC}$ are those described above, with $\equiv^{N_2}_{\langle B,C\rangle} \subseteq \equiv^{N_2'}_{\langle B,C\rangle}$. and $\pi_B(N_2) = \pi_B(N_2')$.

A similar solution is applicable when normalization replaces two-way inclusion dependencies with simple foreign-key dependencies. That example is developed in detail in Examples 3.4.

The main goal of this paper is to place the ideas illustrated by these examples on firm theoretical footing. In contrast to the constant-complement theory, which looks primarily at how a single view may be updated while keeping a second view constant, the focus here is upon how two views may be updated independently. Furthermore, while the work is primarily within the setting of just two views, the long-term goal is nevertheless to address the situation in which there is a larger set of views, as often occurs in the application settings identified above. To this end, the main results are developed without requiring that the views be complementary. Interestingly, complementation does not appear to be a central issue and their is little if any compromise involved.

2 Schemata and Views in a General Framework

Although most of the examples are based upon the relational model, the results of this paper depend only upon the set-theoretic properties of database schemata and views. As such, the underlying framework is basically that employed in the classical papers [6] and [5]. The purpose of this section is to present the essential ideas of that framework in a succinct fashion. The terminology and notation is closest to that employed in [11], to which the reader is referred for details.

Definition 2.1 (Database schemata and morphisms). A database schema $\mathbf{D}$ is modelled completely by its set $\mathsf{LDB}(\mathbf{D})$ of *legal databases* or *states*. A morphism $f : \mathbf{D}_1 \rightarrow \mathbf{D}_2$ of database schemata is represented completely by its underlying function $f : \mathsf{LDB}(\mathbf{D}_1) \rightarrow \mathsf{LDB}(\mathbf{D}_2)$. Since no confusion can result, the morphism and its underlying function will be represented by the same symbol. Of course, schemata may have further structure (such as relational structure), and morphisms may be defined by the relational algebra or calculus, but for this work, it is only the underlying sets and functions which are of formal importance.

Definition 2.2 (Views). A *view* $\Gamma = (\mathbf{V}, \gamma)$ of the schema $\mathbf{D}$ is given by a database schema $\mathbf{V}$ together with a morphism $\gamma : \mathbf{D} \rightarrow \mathbf{V}$ whose underlying function $\gamma : \mathsf{LDB}(\mathbf{D}) \rightarrow \mathsf{LDB}(\mathbf{V})$ is surjective. In a view Γ, the state of its schema $\mathbf{V}$ is always determined completely by the state of the main schema $\mathbf{D}$.

The *congruence* $\mathsf{Congr}(\Gamma)$ of the view Γ is the equivalence relation on $\mathsf{LDB}(\mathbf{D})$ given by $\{(M_1, M_2) \in \mathsf{LDB}(\mathbf{D}) \times \mathsf{LDB}(\mathbf{D}) \mid \gamma(M_1) = \gamma(M_2)\}$. Let $\Gamma_1 = (\mathbf{V}_1, \gamma_1)$

and $\Gamma_2 = (\mathbf{V}_2, \gamma_2)$ be views of the schema $\mathbf{D}$. Write $\Gamma_1 \preceq_{\mathbf{D}} \Gamma_2$ just in case $\mathsf{Congr}(\Gamma_2) \subseteq \mathsf{Congr}(\Gamma_1)$, that is, just in case Γ_2 preserves at least as much information about the state of $\mathbf{D}$ as does Γ_1. The two views Γ_1 and Γ_2 are said to be *isomorphic* if $\mathsf{Congr}(\Gamma_1) = \mathsf{Congr}(\Gamma_2)$; i.e., if $\Gamma_2 \preceq_{\mathbf{D}} \Gamma_1 \preceq_{\mathbf{D}} \Gamma_2$. It is easy to see that $\preceq_{\mathbf{D}}$ is a preorder on the collection of all views of $\mathbf{D}$ and a partial order on the congruences (i.e., on the views up to isomorphism).

A congruence on $\mathsf{LDB}(\mathbf{D})$ may be represented by the partition which it induces [20, Sec. 1]. The partition of $\mathsf{LDB}(\mathbf{D})$ induced by $\mathsf{Congr}(\Gamma)$ is denoted $\mathsf{Partition}(\mathsf{Congr}(\Gamma))$.

Definition 2.3 (Relativized views). Let $\Gamma_1 = (\mathbf{V}_1, \gamma_1)$ and $\Gamma_2 = (\mathbf{V}_2, \gamma_2)$ be views of the schema $\mathbf{D}$. If $\Gamma_1 \preceq_{\mathbf{D}} \Gamma_2$, then Γ_2 may be *relativized* to a view of $\mathbf{V}_1$. More specifically, the function $\lambda\langle \Gamma_1, \Gamma_2 \rangle : \mathsf{LDB}(\mathbf{V}_1) \to \mathsf{LDB}(\mathbf{V}_2)$ is defined via the view congruences by sending a block β of $\mathsf{Partition}(\mathsf{Congr}(\Gamma_1))$ to the block of $\mathsf{Partition}(\mathsf{Congr}(\Gamma_2))$ which contains β. For example, using views of the $\mathbf{E}_0$ introduced in Sec. 1, $\lambda\langle \Pi^{\mathbf{E}_0}_{AB}, \Pi^{\mathbf{E}_0}_{B} \rangle$ sends a state in $\mathsf{LDB}(\Pi^{\mathbf{E}_0}_{AB})$, i.e., a relation for $R_{AB}[AB]$, to its projection on B. In terms of blocks of the equivalence relations, it sends a block β of $\mathsf{Partition}(\mathsf{Congr}(\Gamma))$ consisting of all states with the same projection onto AB, to the block of $\mathsf{Partition}(\mathsf{Congr}(\Pi^{\mathbf{E}_0}_{B}))$ with the projection onto attribute B of the elements of β.

Definition 2.4 (The lattice structure and meets of views). It is a classical result [20, Thm. 5] that the set of all congruences on a set (and hence the set of all views on a database schema) forms a bounded complete lattice (see [9, 2.2 and 2.4] for definitions) under the order induced by $\preceq_{\mathbf{D}}$. More precisely, let Γ_1 and Γ_2 be any views of the schema $\mathbf{D}$. The join $\Gamma_1 \vee \Gamma_2$ is characterized by the congruence $\mathsf{Congr}(\Gamma_1) \cap \mathsf{Congr}(\Gamma_2)$. The join will not be used in this work and so not considered further. More important is the meet $\Gamma_1 \wedge \Gamma_2 = (\mathbf{V}_1 \wedge \mathbf{V}_2, \gamma_1 \wedge \gamma_2)$, which is represented by the intersection of all equivalence relations E on $\mathsf{LDB}(\mathbf{D})$ which satisfy $E \preceq_{\mathbf{D}} \mathsf{Congr}(\Gamma_i)$ for both $i = 1$ and $i = 2$. There is always one such equivalence relation, namely the identity, so the intersection is never over the empty set. An explicit characterization of $\mathsf{Congr}(\Gamma_1 \wedge \Gamma_2)$ may be found in [20, p. 579]. Namely, $(M, M') \in \mathsf{Congr}(\Gamma_1 \wedge \Gamma_2)$ iff there is a chain

$$(M, M_1), (M_1, M_2), \ldots, (M_{i-1}, M_i), (M_i, M_{i+1}), \ldots, (M_{k-1}, M_k), (M_k, M') \tag{cc}$$

of elements in $\mathsf{LDB}(\mathbf{D}) \times \mathsf{LDB}(\mathbf{D})$ in which the right element of a pair matches the left element of its neighbor to the right, and in which each pair is either in $\mathsf{Congr}(\Gamma_1)$ or else in $\mathsf{Congr}(\Gamma_2)$.

While the join of two relational schemata always has a natural representation as a relational schema [14, Def. 3.4], the same cannot be said of the meet. Of course, it always has an abstract representation as a congruence on the states of the main schema, and in many examples, it does have a simple representation. For example, in the context of the schema $\mathbf{E}_0$ of Sec. 1, $\Pi^{\mathbf{E}_0}_{AB} \wedge \Pi^{\mathbf{E}_0}_{BC}$ is represented by the view $\Pi^{\mathbf{E}_0}_{B}$. This is even true for the meet $\Pi^{\mathbf{E}_1}_{AB} \wedge \Pi^{\mathbf{E}_1}_{BC}$ of the views of $\mathbf{E}_1$; this meet is represented by the view $\Pi^{\mathbf{E}_1}_{B}$, the projection onto B.

The greatest view is the identity view, which has the obvious definition and which will not be considered further in this work. The least view is the *zero view*, denoted $\mathsf{ZView}_{\mathbf{D}}$, and has $\mathsf{Congr}(\mathsf{ZView}_{\mathbf{D}}) = \mathsf{LDB}(\mathbf{D}) \times \mathsf{LDB}(\mathbf{D})$. It is a trivial view in that it retains no information about the state of $\mathbf{D}$; its morphism $\mathsf{ZMor}_{\mathbf{D}}$ sends every state of $\mathsf{LDB}(\mathbf{D})$ to the same, single state of the view schema.

Definition 2.5 (Commuting congruences). There is a condition which simplifies the description of the meet given in (cc) of Definition 2.4 above. The pair $\{\Gamma_1 = (\mathbf{V}_1, \gamma_1), \Gamma_2 = (\mathbf{V}_2, \gamma_2)\}$ of views is said to have *commuting congruences* if their if the composition of their congruences is commutative; that is, if $\mathsf{Congr}(\Gamma_1) \circ \mathsf{Congr}(\Gamma_2) = \mathsf{Congr}(\Gamma_2) \circ \mathsf{Congr}(\Gamma_1)$. In this case, the characterization (cc) simplifies considerably. Namely, $(M, M') \in \mathsf{Congr}(\Gamma_1 \wedge \Gamma_2)$ iff there is an $M'' \in \mathsf{LDB}(\mathbf{D})$ such that $(M, M'') \in \mathsf{Congr}(\Gamma_1)$ and $(M'', M') \in \mathsf{Congr}(\Gamma_2)$ (or, equivalently, iff there is an $M'' \in \mathsf{LDB}(\mathbf{D})$ such that $(M, M'') \in \mathsf{Congr}(\Gamma_2)$ and $(M'', M') \in \mathsf{Congr}(\Gamma_1)$) [20, Sec. 8].

Definition 2.6 (Complementary views). The pair $\{\Gamma_1 = (\mathbf{V}_1, \gamma_1), \Gamma_2 = (\mathbf{V}_2, \gamma_2)\}$ of views of $\mathbf{D}$ is called *complementary* if the *decomposition morphism* $\gamma_1 \times \gamma_2 : \mathsf{LDB}(\mathbf{D}) \to \mathsf{LDB}(\mathbf{V}_1) \times \mathsf{LDB}(\mathbf{V}_2)$ given on elements by $M \mapsto (\gamma_1(M), \gamma_2(M))$ is injective. In earlier work, particularly [11], fundamental results were obtained for pairs of views which are both complementary and which have commuting congruences. Such pairs are called *meet complementary.* In this work, the property of being complementary will not be of central importance, but it will still be mentioned in some discussion of the results.

Definition 2.7 (Updates and Reflections). An *update* on the schema $\mathbf{D}$ is just a pair $(M_1, M_2) \in \mathsf{LDB}(\mathbf{D}) \times \mathsf{LDB}(\mathbf{D})$. Think of M_1 as the state before the update operation and M_2 as the state afterwards. The set of all updates on $\mathbf{D}$ is denoted $\mathsf{Updates}(\mathbf{D})$.

Given a view $\Gamma = (\mathbf{V}, \gamma)$ of $\mathbf{D}$ and an update $u = (N_1, N_2) \in \mathsf{Updates}(\mathbf{V})$, a *reflection* (or *translation*) *of u along Γ* is a $u' = (M_1, M_2) \in \mathsf{Updates}(\mathbf{D})$ with the property that $\gamma(M_i) = N_i$ for $i \in \{1, 2\}$. In this case, u' is also called a reflection (or translation) of u *for* M_1 along Γ. The set of all reflections of u along Γ is denoted $\mathsf{Reflections}_{\Gamma}\langle u \rangle$.

3 Basic Theory of Independent Update Strategies

In this section, the central ideas surrounding independent update strategies are developed. Some of these, particularly those involving commuting congruences, have already been developed in part in the context of complementary pairs [11]. However, the focus here is not at all upon complements. Indeed, the assumption that the views under consideration are complementary is never made. Furthermore, while the emphasis in [11] is upon the constant-complement update strategy in the presence of meet complements, the main focus here is upon situations in which the meet property (i.e., commuting congruences) fails to hold. This presentation is independent of [11], and does not require knowledge of the specific results of that paper.

Notation 3.1 (Running schema and views). Throughout this section, unless stated specifically to the contrary, take $\mathbf{D}$ to be a database schema and $\Gamma_1 = (\mathbf{V}_1, \gamma_1)$ and $\Gamma_2 = (\mathbf{V}_2, \gamma_2)$ to be views of $\mathbf{D}$. Γ_1 and Γ_2 need not be complements of each other.

Definition 3.2 (Updates relative to a second view). The goal is to identify properties on subsets of $\mathsf{Updates}(\mathbf{V}_1)$ which characterize useful yet independent update strategies. To this end, there are three distinct notions of independence which are of importance. In that which follows, let $u = (N, N') \in \mathsf{Updates}(\mathbf{V}_1)$, and define $\mathsf{Reflections}_{\Gamma_1|\Gamma_2}\langle u\rangle$ to be the subset of $\mathsf{Reflections}_{\Gamma_1}\langle u\rangle$ which keeps the state of Γ_2 constant. More precisely, $\mathsf{Reflections}_{\Gamma_1|\Gamma_2}\langle u\rangle = \{(M_1, M_2) \in \mathsf{Reflections}_{\Gamma_1}\langle u\rangle \mid (M_1, M_2) \in \mathsf{Congr}(\Gamma_2)\}$.

(a) Call u *somewhere Γ_2-independent* if for some $M_1 \in \mathsf{LDB}(\mathbf{D})$ with $\gamma_1(M_1) = N$, there is an $M_2 \in \mathsf{LDB}(\mathbf{D})$ with the property that $(M_1, M_2) \in \mathsf{Reflections}_{\Gamma_1|\Gamma_2}\langle u\rangle$. The set of all somewhere Γ_2-independent updates on Γ_1 is denoted $\mathsf{IndUpd}_{\exists}\langle \Gamma_1|\Gamma_2\rangle$.

Thus, u is somewhere Γ_2-independent if the update may be made for some states of the view Γ_2, but not necessarily all. The update $(\{R_{AB}(\mathrm{a}_1, \mathrm{b}_1), R_{AB}(\mathrm{a}_2, \mathrm{b}_2)\}, \{R_{AB}(\mathrm{a}_1, \mathrm{b}_1), R_{AB}(\mathrm{a}_1, \mathrm{b}_2)\})$ on $\Pi^{\mathbf{E}_1}_{AB}$ of Sec. 1 is an example which is somewhere $\Pi^{\mathbf{E}_1}_{BC}$-independent. It is not, however, everywhere independent, since there states of the view $\Pi^{\mathbf{E}_1}_{BC}$, such as $\{R_{BC}(\mathrm{b}_1, \mathrm{c}_1), R_{BC}(\mathrm{b}_2, \mathrm{c}_2)\}$, for which it cannot be realized without changing that state.

(b) Call u *everywhere Γ_2-independent* if for every $M_1 \in \mathsf{LDB}(\mathbf{D})$ with $\gamma_1(M_1) = N$, there is an $M_2 \in \mathsf{LDB}(\mathbf{D})$ with the property that $(M_1, M_2) \in \mathsf{Reflections}_{\Gamma_1|\Gamma_2}\langle u\rangle$. The set of all everywhere Γ_2-indepdendent updates on Γ_1 is denoted $\mathsf{IndUpd}_{\forall}\langle \Gamma_1|\Gamma_2\rangle)$.

The update $(\{R_{AB}(\mathrm{a}_1, \mathrm{b}_1), R_{AB}(\mathrm{a}_2, \mathrm{b}_2)\}, \{R_{AB}(\mathrm{a}_1, \mathrm{b}_1), R_{AB}(\mathrm{a}_3, \mathrm{b}_2)\})$ on $\Pi^{\mathbf{E}_1}_{AB}$ is an example which is everywhere $\Pi^{\mathbf{E}_1}_{BC}$-independent.

The third notion of independence characterizes independence in the situation when the congruences of Γ_1 and Γ_2 commute, and so is closely tied to the theory of constant-complement updates as presented in [11].

(c) Call u *meetwise Γ_2-independent* if $\lambda\langle \Gamma_1, \Gamma_1 \wedge \Gamma_2\rangle(N) = \lambda\langle \Gamma_1, \Gamma_1 \wedge \Gamma_2\rangle(N')$. The set of all meetwise Γ_2-indepdendent updates on Γ_1 is denoted $\mathsf{IndUpd}_{\wedge}\langle \Gamma_1|\Gamma_2\rangle$.

Given that $\Pi^{\mathbf{E}_1}_{AB} \wedge \Pi^{\mathbf{E}_1}_{BC} = \Pi^{\mathbf{E}_1}_{B}$, the update $(\{R_{AB}(\mathrm{a}_1, \mathrm{b}_1), R_{AB}(\mathrm{a}_2, \mathrm{b}_2)\}, \{R_{AB}(\mathrm{a}_1, \mathrm{b}_1), R_{AB}(\mathrm{a}_3, \mathrm{b}_2)\})$ on $\Pi^{\mathbf{E}_1}_{AB}$ is meetwise $\Pi^{\mathbf{E}_1}_{BC}$-independent since the projection onto $\Pi^{\mathbf{E}_1}_{B}$ is $\{R_B(\mathrm{b}_1), R_B(\mathrm{b}_2)\}$ in each case.

For each of these three notions, there is a corresponding definition of reflected updates. Specifically, define

$$\begin{aligned}
\mathsf{ReflIndUpd}_{\exists}\langle \Gamma_1|\Gamma_2\rangle &= \{\mathsf{Reflections}_{\Gamma_1|\Gamma_2}\langle u\rangle \mid u \in \mathsf{IndUpd}_{\exists}\langle \Gamma_1|\Gamma_2\rangle\},\\
\mathsf{ReflIndUpd}_{\forall}\langle \Gamma_1|\Gamma_2\rangle &= \{\mathsf{Reflections}_{\Gamma_1|\Gamma_2}\langle u\rangle \mid u \in \mathsf{IndUpd}_{\forall}\langle \Gamma_1|\Gamma_2\rangle\},\\
\text{and } \mathsf{ReflIndUpd}_{\wedge}\langle \Gamma_1|\Gamma_2\rangle &= \{\mathsf{Reflections}_{\Gamma_1|\Gamma_2}\langle u\rangle \mid u \in \mathsf{IndUpd}_{\wedge}\langle \Gamma_1|\Gamma_2\rangle\}.
\end{aligned}$$

In these definitions, there is no assumption that Γ_1 and Γ_2 be complements. However, if they are complements, then each of $\mathsf{ReflIndUpd}_{\exists}\langle \Gamma_1|\Gamma_2\rangle$, $\mathsf{ReflIndUpd}_{\forall}\langle \Gamma_1|\Gamma_2\rangle$, and $\mathsf{ReflIndUpd}_{\wedge}\langle \Gamma_1|\Gamma_2\rangle$ must be a function, in the sense that if $(M_1, M_2), (M_1, M_2') \in \mathsf{Updates}(\mathbf{D})$ are both in any of these sets, then $M_2 = M_2'$. If Γ_1 and Γ_2 are complements, this means that the decomposition morphism $\gamma_1 \times \gamma_2$ (Definition 2.6) must be injective, and so M_2 must be the unique element $(\gamma_1 \times \gamma_2)^{-1}(N_2, \gamma_2(M_2))$, if it exists.

All three notions of Γ_2-independence recapture locality of effect, as defined in Sec. 1. While only everywhere Γ_2 independence recaptures context independence, the other two provide crucial insights into what can go wrong and how things can be extended. As a first step, the question of whether or not these are equivalence relations is examined.

Observation 3.3 (Reflexivity and transitivity).

(a) *Each of* $\mathsf{IndUpd}_{\wedge}\langle \Gamma_1 | \Gamma_2 \rangle$, $\mathsf{IndUpd}_{\exists}\langle \Gamma_1 | \Gamma_2 \rangle$, $\mathsf{ReflIndUpd}_{\wedge}\langle \Gamma_1 | \Gamma_2 \rangle$, *and* $\mathsf{ReflIndUpd}_{\exists}\langle \Gamma_1 | \Gamma_2 \rangle$ *is an equivalence relations.*

(b) $\mathsf{IndUpd}_{\forall}\langle \Gamma_1 | \Gamma_2 \rangle$ *and* $\mathsf{ReflIndUpd}_{\forall}\langle \Gamma_1 | \Gamma_2 \rangle$ *are reflexive and transitive, but not necessarily symmetric. Thus, they need not be equivalence relations.*

Proof. All of the "positive" conditions are routine verifications, which are left to the reader. That $\mathsf{IndUpd}_{\forall}\langle \Gamma_1 | \Gamma_2 \rangle$ and $\mathsf{ReflIndUpd}_{\forall}\langle \Gamma_1 | \Gamma_2 \rangle$ need not be symmetric is illustrated in Examples 3.4, immediately below. □

Examples 3.4 (Non-reversible independent updates). To illustrate the idea of non-reversible updates, consider the schema $\mathbf{E}_2$ with two relation symbols $R_{AB}[AB]$ and $R_{BC}[BC]$. The latter relation is governed by the FD $B \to C$, and, in addition, the two relations are connected via the foreign-key dependency $R_{AB}[B] \subseteq R_{BC}[B]$. Define $\Pi_{AB}^{\mathbf{E}_2} = (\mathbf{E}_2^{AB}, \pi_{AB}^{\mathbf{E}_2})$ and $\Pi_{BC}^{\mathbf{E}_2} = (\mathbf{E}_2^{AB}, \pi_{AB}^{\mathbf{E}_2})$ as the views which preserve $R_{AB}[AB]$ and $R_{BC}[BC]$, respectively, and for $N \in \mathsf{LDB}(\mathbf{E}_2^{AB})$ or $N \in \mathsf{LDB}(\mathbf{E}_2^{BC})$, let $\pi_B(N)$ denote $\{\mathrm{b} \mid (\exists t \in N)(t[B] = b\}$.

It is easy to see that $\mathsf{IndUpd}_{\forall}\langle \Pi_{AB}^{\mathbf{E}_3} | \Pi_{BC}^{\mathbf{E}_2} \rangle$ is the set of all updates $(N, N') \in \mathsf{Updates}(\mathbf{V}_1)$ for which $\pi_B(N') \subseteq \pi_B(N)$. A tuple of the form $R_{AB}(\mathrm{a}, \mathrm{b})$ may always be deleted, even if there is no other tuple of the form $R_{AB}(x, \mathrm{b})$, but a tuple of the form $R_{AB}(\mathrm{a}, \mathrm{b})$ may not be added if there is not already another of the form $R_{AB}(x, \mathrm{b})$. Thus, if $R_{AB}(\mathrm{a}, \mathrm{b})$ is deleted, it may not be reinserted. For $\mathsf{IndUpd}_{\forall}\langle \Pi_{AB}^{\mathbf{E}_3} | \Pi_{BC}^{\mathbf{E}_2} \rangle$, the situation is reversed; $(N, N') \in \mathsf{IndUpd}_{\forall}\langle \Pi_{BC}^{\mathbf{E}_2} | \Pi_{AB}^{\mathbf{E}_2} \rangle$ iff $\pi_B(N) \subseteq \pi_B(N')$. A tuple of the form $R_{BC}(\mathrm{b}, \mathrm{c})$ may always be inserted, but not deleted unless there is another tuple of the form $R_{BC}(\mathrm{b}, x)$. Hence, neither $\mathsf{IndUpd}_{\forall}\langle \Pi_{AB}^{\mathbf{E}_3} | \Pi_{BC}^{\mathbf{E}_2} \rangle$ nor $\mathsf{IndUpd}_{\forall}\langle \Pi_{BC}^{\mathbf{E}_3} | \Pi_{AB}^{\mathbf{E}_2} \rangle$ is symmetric.

A similar situation governs the example surrounding $\mathbf{E}_1$ of Sec. 1. A view update $(N, N') \in \mathsf{IndUpd}_{\forall}\langle \Pi_{AB}^{\mathbf{E}_1} | \Pi_{BC}^{\mathbf{E}_1} \rangle$ is allowed if $\pi_B(N) = \pi_B(N')$ and $\equiv_{\langle B,A \rangle}^{N'} \subseteq \equiv_{\langle B,A \rangle}^{N}$, but not if $\equiv_{\langle B,A \rangle}^{N'} \subsetneq \equiv_{\langle B,A \rangle}^{N}$. Thus, a view update of the form $(\{R_{AB}(\mathrm{a}_1, \mathrm{b}_1), R_{AB}(\mathrm{a}_2, \mathrm{b}_2)\}, \{R_{AB}(\mathrm{a}_1, \mathrm{b}_1), R_{AB}(\mathrm{a}_1, \mathrm{b}_2)\})$ is not allowed. An analogous condition holds for $\mathsf{IndUpd}_{\forall}\langle \Pi_{BC}^{\mathbf{E}_1} | \Pi_{AB}^{\mathbf{E}_1} \rangle$.

These examples suggest the way to extend the notion of independent update. Return to $\mathbf{E}_2$ and its views. Suppose that the state of the schema $\mathbf{E}_2$ is $M_{21} = \{R_{AB}(\mathrm{a}_1, \mathrm{b}_1), R_{AB}(\mathrm{a}_2, \mathrm{b}_2), R_{BC}(\mathrm{b}_1, \mathrm{c}_1), R_{BC}(\mathrm{b}_2, \mathrm{c}_2)\}$. Then the update $u_{21} = (\{R_{AB}(\mathrm{a}_1, \mathrm{b}_1), R_{AB}(\mathrm{a}_2, \mathrm{b}_2)\}, \{R_{AB}(\mathrm{a}_1, \mathrm{b}_1)\})$ on $\Pi_{AB}^{\mathbf{E}_2}$ is in $\mathsf{IndUpd}_{\forall}\langle \Pi_{AB}^{\mathbf{E}_2} | \Pi_{BC}^{\mathbf{E}_2} \rangle$ but the reverse update $u_{21}' = (\{R_{AB}(\mathrm{a}_1, \mathrm{b}_1)\}, \{R_{AB}(\mathrm{a}_1, \mathrm{b}_1), R_{AB}(\mathrm{a}_2, \mathrm{b}_2)\})$ is not. However, if it is known that the state of $\Pi^{\mathbf{E}_2}_{BC}$ did not change after the execution of u_{21}, then a subsequent execution of u'_{21} is indeed possible while keeping $\Pi^{\mathbf{E}_2}_{BC}$ constant. Even stronger, if updates in $\mathsf{IndUpd}_{\forall}\langle \Pi^{\mathbf{E}_2}_{BC} | \Pi^{\mathbf{E}_2}_{AB}\rangle$ are applied, this condition nevertheless continues to hold. It is only if an update to $\Pi^{\mathbf{E}_2}_{BC}$ is applied which removes elements from $\pi_B(\{R_{AB}(\mathrm{a}_1,\mathrm{b}_1), R_{AB}(\mathrm{a}_2,\mathrm{b}_2)\}) = \{\mathrm{b}_1,\mathrm{b}_2\}$ that the update u_{21} may become irreversible. The strategy is that non-reversible updates in $\mathsf{IndUpd}_{\forall}\langle \Gamma_1|\Gamma_2\rangle$ may in fact be reversed provided that the only updates to Γ_2 which are allowed are those in $\mathsf{IndUpd}_{\forall}\langle \Gamma_2|\Gamma_1\rangle$ and their reversals, all within the context of a given initial state. A systematic development of these ideas constitutes the remainder of this paper.

Proposition 3.5 (Comparison of the three notions of independent updates). $\mathsf{IndUpd}_{\forall}\langle \Gamma_1|\Gamma_2\rangle \subseteq \mathsf{IndUpd}_{\exists}\langle \Gamma_1|\Gamma_2\rangle \subseteq \mathsf{IndUpd}_{\wedge}\langle \Gamma_1|\Gamma_2\rangle$. *Furthermore, their exist examples for which these inclusions are proper.*

Proof. That $\mathsf{IndUpd}_{\forall}\langle \Gamma_1|\Gamma_2\rangle \subseteq \mathsf{IndUpd}_{\exists}\langle \Gamma_1|\Gamma_2\rangle$ is immediate, and an example for which the inclusion is proper is given by $\mathbf{E}_1$ and its views in Sec. 1.

It is also easy to see that $\mathsf{IndUpd}_{\exists}\langle \Gamma_1|\Gamma_2\rangle \subseteq \mathsf{IndUpd}_{\wedge}\langle \Gamma_1|\Gamma_2\rangle$, since if $(N,N') \in \mathsf{IndUpd}_{\exists}\langle \Gamma_1|\Gamma_2\rangle$, then by definition there is a pair $(M,M') \in \mathsf{Reflections}_{\Gamma_1|\Gamma_2}\langle (N,N')\rangle$, and $(M,M') \in \mathsf{Congr}(\Gamma_2)$, since the update holds Γ_2 constant. Then, since $\mathsf{Congr}(\Gamma_2) \subseteq \mathsf{Congr}(\Gamma_1 \wedge \Gamma_2)$, the results follows. For an example in which this inclusion is proper, let $\mathbf{E}_3$ have five states: $\mathsf{LDB}(\mathbf{E}_3) = \{a,b,c,d,e\}$. Let $\Omega_{31} = (\mathbf{V}_{31},\omega_{31})$ be the view with $\mathsf{Partition}(\mathsf{Congr}(\Omega_{31})) = \{\{a,b\},\{c,d\},\{e\}\}$ and let $\Omega_{32} = (\mathbf{V}_{32},\omega_{32})$ have $\mathsf{Partition}(\mathsf{Congr}(\Omega_{32})) = \{\{a\},\{b,c\},\{d,e\}\}$. It is easy to show that $\mathsf{Partition}(\mathsf{Congr}(\omega_{31} \wedge \omega_{32})) = \{a,b,c,d,e\}$, i.e., it is defined by $\mathsf{ZView}_{\mathbf{E}_3}$. This means that $\mathsf{IndUpd}_{\wedge}\langle \Omega_{31}|\Omega_{32}\rangle = \mathsf{LDB}(\mathbf{E}_3)\times\mathsf{LDB}(\mathbf{E}_3)$; i.e., any update is allowed. However, the update $(\{a,b\},\{e\})$ is not in $\mathsf{IndUpd}_{\exists}\langle \Omega_{31}|\Omega_{32}\rangle$, and so $\mathsf{IndUpd}_{\wedge}\langle \Omega_{31}|\Omega_{32}\rangle$ is a proper subset of it. □

Definition 3.6 (Compatible pairs and independent update pairs). A *compatible pair* for $\{\Gamma_1,\Gamma_2\}$ is an $(N_1,N_2) \in \mathsf{LDB}(\mathbf{V}_1) \times \mathsf{LDB}(\mathbf{V}_2)$ which arises from some $M \in \mathsf{LDB}(\mathbf{D})$. If $\{\Gamma_1,\Gamma_2\}$ forms a complementary pair, then there is at most one compatible pair associated with each $M \in \mathsf{LDB}(\mathbf{D})$; namely $(\gamma_1 \wedge \gamma_2)^{-1}(N_1,N_2)$ when it exists. However, in the more general context, there may be many, since $\gamma_1 \times \gamma_2$ need not be injective. Formally, define $\mathsf{Compat}\langle \Gamma_1;\Gamma_2\rangle = \{(N_1,N_2) \in \mathsf{LDB}(\mathbf{V}_1) \times \mathsf{LDB}(\mathbf{V}_2) \mid (\exists M \in \mathsf{LDB}(\mathbf{D}))(\forall i \in \{1,2\})(\gamma_i(M) = N_i)\}$.

An *independent update pair* is a pair of updates $(u_1,u_2) \in \mathsf{Updates}(\mathbf{V}_1) \times \mathsf{Updates}(\mathbf{V}_2)$ which may be executed independently of one another. Formally, define $\mathsf{IndUpd}_{\forall}\langle \Gamma_1\|\Gamma_2\rangle = \{(N_1,N_2),(N'_1,N'_2) \mid (N_1,N_2) \in \mathsf{Compat}\langle \Gamma_1;\Gamma_2\rangle$ and $(N_1,N'_1) \in \mathsf{IndUpd}_{\forall}\langle \Gamma_1|\Gamma_2\rangle$ and $(N_2,N'_2) \in \mathsf{IndUpd}_{\forall}\langle \Gamma_2|\Gamma_1\rangle\}$.

Updates in $\mathsf{IndUpd}_{\forall}\langle \Gamma_1\|\Gamma_2\rangle$ may be performed individually as updates in $\mathsf{IndUpd}_{\forall}\langle \Gamma_1|\Gamma_2\rangle$ and $\mathsf{IndUpd}_{\forall}\langle \Gamma_2|\Gamma_1\rangle$, as well as concurrently, and these operations all preserve compatibility. Formally, this is expressed as follows.

Proposition 3.7 (Independent updates). *Let* $(N_1,N_2) \in \mathsf{Compat}\langle \Gamma_1;\Gamma_2\rangle$ *and let* $((N_1,N'_1),(N_2,N'_2)) \in \mathsf{IndUpd}_{\forall}\langle \Gamma_1\|\Gamma_2\rangle$. *Then each of the pairs*

$((N_1, N_1'), (N_2, N_2))$, $((N_1, N_1'), (N_2', N_2'))$, $((N_1, N_1), (N_2, N_2'))$, *and* $((N_1', N_1'), (N_2, N_2'))$ *is in* $\mathsf{IndUpd}_\forall\langle\Gamma_1\|\Gamma_2\rangle$ *as well. In particular, each of* (N_1', N_2'), (N_1', N_2), *and* (N_1', N_2') *is in* $\mathsf{Compat}\langle\Gamma_1;\Gamma_2\rangle$.

Proof. It suffices to equate certain elements in the description of Definition 3.6. For example, letting N_2' be N_2, which is always possible since identity updates such as (N_2, N_2) are in $\mathsf{IndUpd}_\forall\langle\Gamma_2|\Gamma_1\rangle$ regardless of the choices of Γ_1 and Γ_2, it follows that $((N_1, N_1'), (N_2, N_2)) \in \mathsf{IndUpd}_\forall\langle\Gamma_1\|\Gamma_2\rangle$. The other three cases are shown similarly. □

The next theorem provides a comprehensive characterization of the conditions for independent updates, without any requirement of complementation. The equivalence of (a), (b), and (c) has already been shown in [11, 2.14] for the special case of complementary pairs, using a different approach [11, Thm. 2.14].

Theorem 3.8 (Independence and commuting congruences). *The following conditions are equivalent.*

(a) *The pair* $\{\Gamma_1, \Gamma_2\}$ *has commuting congruences; i.e.,* $\mathsf{Congr}(\Gamma_1)\circ\mathsf{Congr}(\Gamma_2) = \mathsf{Congr}(\Gamma_2)\circ\mathsf{Congr}(\Gamma_1)$.

(b) *For any* $N_1 \in \mathsf{LDB}(\mathbf{V}_1)$ *and* $N_2 \in \mathsf{LDB}(\mathbf{V}_2)$, $(N_1, N_2) \in \mathsf{Compat}\langle\Gamma_1;\Gamma_2\rangle$ *iff* $\lambda\langle\Gamma_1, \Gamma_1 \wedge \Gamma_2\rangle(N_1) = \lambda\langle\Gamma_2, \Gamma_1 \wedge \Gamma_2\rangle(N_2)$.

(c) *For any* $N_1, N_1' \in \mathsf{LDB}(\mathbf{V}_1)$ *and* $N_2, N_2' \in \mathsf{LDB}(\mathbf{V}_2)$, *if any three elements of the set* $\{(N_1, N_2), (N_1, N_2'), (N_1', N_2), (N_1', N_2')\}$ *are in* $\mathsf{Compat}\langle\Gamma_1;\Gamma_2\rangle$, *then so too is the fourth.*

(d) $\mathsf{IndUpd}_\forall\langle\Gamma_1|\Gamma_2\rangle = \mathsf{IndUpd}_\wedge\langle\Gamma_1|\Gamma_2\rangle$.

(e) $\mathsf{IndUpd}_\forall\langle\Gamma_1|\Gamma_2\rangle = \mathsf{IndUpd}_\exists\langle\Gamma_1|\Gamma_2\rangle$.

(f) $\mathsf{ReflIndUpd}_\forall\langle\Gamma_1|\Gamma_2\rangle = \mathsf{Congr}(\Gamma_2)$.

(g) $\mathsf{ReflIndUpd}_\forall\langle\Gamma_1|\Gamma_2\rangle$ *is an equivalence relation.*

(h) $\mathsf{IndUpd}_\forall\langle\Gamma_1|\Gamma_2\rangle$ *is an equivalence relation.*

Proof. ((a) ⇒ (b)): First, assume that $\lambda\langle\Gamma_1, \Gamma_1 \wedge \Gamma_2\rangle(N_1) = \lambda\langle\Gamma_2, \Gamma_1 \wedge \Gamma_2\rangle(N_2)$, and let $M_1, M_2 \in \mathsf{LDB}(\mathbf{D})$ with $\gamma_1(M_1) = N_1$ and $\gamma_2(M_2) = N_2$. Then $(\gamma_1 \wedge \gamma_2)(M_1) = (\gamma_1 \wedge \gamma_2)(M_2)$, i.e., $(M_1, M_2) \in \mathsf{Congr}(\Gamma_1 \wedge \Gamma_2)$. Using the characterization of $\mathsf{Congr}(\Gamma_1 \wedge \Gamma_2)$ for commuting congruences given in Definition 2.4, there must be an $M \in \mathsf{LDB}(\mathbf{D})$ with $(M_1, M) \in \mathsf{Congr}(\Gamma_1)$ and $(M, M_2) \in \mathsf{Congr}(\Gamma_2)$. Furthermore, $\gamma_1(M) = \gamma_1(M_1) = N_1$, and $\gamma_2(M) = \gamma_1(M_2) = N_2$, whence $(N_1, N_2) \in \mathsf{Compat}\langle\Gamma_1;\Gamma_2\rangle$. Conversely, if $(N_1, N_2) \in \mathsf{Compat}\langle\Gamma_1;\Gamma_2\rangle$, then there exists an $M \in \mathsf{LDB}(\mathbf{D})$ with $\gamma_1(M) = N_1$ and $\gamma_2(M) = N_2$. Since this M maps to a single block of $\mathsf{Partition}(\mathsf{Congr}(\Gamma_1 \wedge \Gamma_2))$, N_1 and N_2 must be associated with the same block as well.

((b) ⇒ (c)): Immediate.

((c) ⇒ (e)): Let $(N_1, N_1') \in \mathsf{IndUpd}_\exists\langle\Gamma_1|\Gamma_2\rangle$, and choose $(M_1, M_2) \in \mathsf{Reflections}_{\Gamma_1|\Gamma_2}\langle(N_1, N_2)\rangle$. Then $(N_1, \gamma_2(M_1)), (N_1', \gamma_2(M_1)) \in \mathsf{Compat}\langle\Gamma_1;\Gamma_2\rangle$. Choose $M_1' \in \mathsf{LDB}(\mathbf{D})$ with $\gamma_1(M_1') = N_1$. Then $(N_1, \gamma_2(M_1')) \in \mathsf{Compat}\langle\Gamma_1;\Gamma_2\rangle$ as well. Hence, by (c), $(N_1', \gamma_2(M_1')) \in \mathsf{Compat}\langle\Gamma_1;\Gamma_2\rangle$, whence $(N_1, N_1') \in \mathsf{IndUpd}_\forall\langle\Gamma_1|\Gamma_2\rangle$, as required.

((e) $\Rightarrow$ (f)): It is immediate that $\mathsf{ReflIndUpd}_{\forall}\langle\Gamma_1|\Gamma_2\rangle \subseteq \mathsf{Congr}(\Gamma_2)$. Conversely, let $(M_1, M_2) \in \mathsf{Congr}(\Gamma_2)$. Then $(\gamma_1(M_1), \gamma_1(M_2)) \in \mathsf{IndUpd}_{\exists}\langle\Gamma_1|\Gamma_2\rangle$, just by construction. Hence, invoking (e), $(\gamma_1(M_1), \gamma_1(M_2)) \in \mathsf{IndUpd}_{\forall}\langle\Gamma_1|\Gamma_2\rangle$ as well, whence $(M_1, M_2) \in \mathsf{ReflIndUpd}_{\forall}\langle\Gamma_1|\Gamma_2\rangle$ and so $\mathsf{ReflIndUpd}_{\exists}\langle\Gamma_1|\Gamma_2\rangle = \mathsf{ReflIndUpd}_{\forall}\langle\Gamma_1|\Gamma_2\rangle$.

((f) $\Rightarrow$ (g)): Immediate.

((g) $\Rightarrow$ (h)): The proof is a routine verification.

((h) $\Rightarrow$ (a)): Let $(M_1, M_2) \in \mathsf{Congr}(\Gamma_1) \circ \mathsf{Congr}(\Gamma_2)$. Then there is an $M' \in \mathsf{LDB}(\mathbf{D})$ with $(M_1, M') \in \mathsf{Congr}(\Gamma_1)$ and $(M', M_2) \in \mathsf{Congr}(\Gamma_2)$. Since $\gamma_2(M') = \gamma_2(M_2)$ and $(M_1, M') \in \mathsf{IndUpd}_{\exists}\langle\Gamma_1|\Gamma_2\rangle$, it follows that $(\gamma_1(M_1), \gamma_1(M_2)) \in \mathsf{IndUpd}_{\exists}\langle\Gamma_1|\Gamma_2\rangle$. Now, choose any $M_1' \in \mathsf{LDB}(\mathbf{D})$ with $\gamma_1(M_1') = \gamma_1(M_1)$. Then $(M_1', M_2) = (M_1', M_1) \circ (M_1, M_2) \in \mathsf{Congr}(\Gamma_1) \circ \mathsf{Congr}(\Gamma_1) \circ \mathsf{Congr}(\Gamma_2) = \mathsf{Congr}(\Gamma_1) \circ \mathsf{Congr}(\Gamma_2)$, and so $(\gamma_1(M_1), \gamma_1(M_2)) \in \mathsf{IndUpd}_{\exists}\langle\Gamma_1|\Gamma_2\rangle$ as well. Since M_1' was arbitrary with $\gamma(M_1) = \gamma(M_1')$, it follows that $(\gamma_1(M_1), \gamma_1(M_2)) \in \mathsf{IndUpd}_{\forall}\langle\Gamma_1|\Gamma_2\rangle$. Conversely, if $(M_1, M_2) \in \mathsf{Updates}(\mathbf{D})$ with $(\gamma_1(M_1), \gamma_1(M_2)) \in \mathsf{IndUpd}_{\forall}\langle\Gamma_1|\Gamma_2\rangle$, then there must be an $M' \in \mathsf{LDB}(\mathbf{D})$ with $(M_1, M') \in \mathsf{Congr}(\Gamma_1)$ and $(M', M_2) \in \mathsf{Congr}(\Gamma_2)$; i.e., $(M'M_2) \in \mathsf{ReflIndUpd}_{\forall}\langle\Gamma_1|\Gamma_2\rangle$. In other words, $(M_1, M_2) \in \mathsf{Congr}(\Gamma_1) \circ \mathsf{Congr}(\Gamma_2)$. Thus, $(M_1, M_2) \in \mathsf{Congr}(\Gamma_1) \circ \mathsf{Congr}(\Gamma_2)$ iff $(\gamma_1(M_1), \gamma_2(M_2)) \in \mathsf{IndUpd}_{\forall}\langle\Gamma_1|\Gamma_2\rangle$. Since $\mathsf{IndUpd}_{\forall}\langle\Gamma_1|\Gamma_2\rangle$ is assumed to be an equivalence relation, it follows easily that $\mathsf{Congr}(\Gamma_1) \circ \mathsf{Congr}(\Gamma_2)$ must be an equivalence relation as well. Then $\mathsf{Congr}(\Gamma_1)\circ\mathsf{Congr}(\Gamma_2) = \mathsf{Congr}(\Gamma_2)\circ\mathsf{Congr}(\Gamma_1)$ follows immediately, since one is the reverse of the other; i.e., $(M_1, M_2) \in \mathsf{Congr}(\Gamma_1) \circ \mathsf{Congr}(\Gamma_2)$ iff $(M_2, M_1) \in \mathsf{Congr}(\Gamma_2) \circ \mathsf{Congr}(\Gamma_1)$.

((a) $\Rightarrow$ (d)): Let $(N_1, N_1') \in \mathsf{IndUpd}_{\wedge}\langle\Gamma_1|\Gamma_2\rangle$. Then, as described in Definition 2.4, for any $(M, M') \in \mathsf{Reflections}_{\Gamma_1|\Gamma_2}\langle(N, N')\rangle$, there is a sequence $M_1, M_2, \ldots, M_k$ of elements of $\mathsf{LDB}(\mathbf{D})$ with the property that $M_1 = M$, $M_k = M'$, for each odd i, $1 \le i \le k$, $(M_i, M_{i+1}) \in \mathsf{Congr}(\Gamma_2)$ and for each even i, $1 \le i \le k$, $(M_i, M_{i+1}) \in \mathsf{Congr}(\Gamma_1)$. However, in view of condition (a), which guarantees commuting congruences, it follows also from the discussion of Definition 2.4 that k may be chosen to be 3. That is, there are $(M_1, M_2) \in \mathsf{Congr}(\Gamma_2)$ and $(M_2, M_3) \in \mathsf{Congr}(\Gamma_1)$ with $\gamma_1(M_1) = N_1$, $\gamma_1(M_3) = N_1'$, and $\gamma_2(M_2) = \gamma_2(M_3)$. Since M may be chosen arbitrarily with the property that $\gamma_1(M_1) = N$, this means in particular that $(N_1, N_1') = (\gamma_1(M_1), \gamma_1(M_2)) \in \mathsf{IndUpd}_{\forall}\langle\Gamma_1|\Gamma_2\rangle$, as required. (That M may be chosen arbitrarily follows from the fact that the equivalence relation $\mathsf{Congr}(\Gamma_1 \wedge \Gamma_2)$ is coarser than $\mathsf{Congr}(\Gamma_2)$, and so any two elements of $\mathsf{LDB}(\mathbf{D})$ which are equivalent under $\mathsf{Congr}(\Gamma_2)$ are equivalent under $\mathsf{Congr}(\Gamma_1) \wedge \mathsf{Congr}(\Gamma_2)$ as well.)

((d) $\Rightarrow$ (e)): This follows immediately from Proposition 3.5. □

The thrust of this result is that as soon as $\mathsf{IndUpd}_{\forall}\langle\Gamma_1|\Gamma_2\rangle$ becomes an equivalence relation, then the classical characterization in terms of commuting congruences and meet dependencies (a)-(c) takes hold, and each of the concepts of independent update $\mathsf{IndUpd}_{\forall}\langle\Gamma_1|\Gamma_2\rangle$, $\mathsf{IndUpd}_{\exists}\langle\Gamma_1|\Gamma_2\rangle$, and $\mathsf{IndUpd}_{\wedge}\langle\Gamma_1|\Gamma_2\rangle$ becomes equivalent to all of the others. There is furthermore a symmetry in results (d)-(f); if Γ_1 and Γ_2 are swapped in any or all of these, the result remains valid.

In particular, $\mathsf{IndUpd}_\forall\langle \Gamma_1|\Gamma_2\rangle$ is an equivalence relation iff $\mathsf{IndUpd}_\forall\langle \Gamma_2|\Gamma_1\rangle$ is. In other words, if independent updates are well behaved on Γ_1, then they are well behaved on Γ_2 as well.

The question becomes, then, how to recapture the extended updates identified in the examples of Sec. 1 and Examples 3.4. The answer is that rather than trying to avoid allowing $\mathsf{IndUpd}_\forall\langle \Gamma_1|\Gamma_2\rangle$ to become an equivalence relation (which in view of the above result would imply many other limitations), the set of allowable legal databases is trimmed so that $\mathsf{IndUpd}_\forall\langle \Gamma_1|\Gamma_2\rangle$ (and so $\mathsf{IndUpd}_\forall\langle \Gamma_2|\Gamma_1\rangle$ as well) becomes an equivalence relation on that which remains. The key idea is to start with a pair $(N_1, N_2) \in \mathsf{LDB}(\mathbf{V}_1)\times\mathsf{LDB}(\mathbf{V}_2)$, and then restrict attention to those states which can be reached from those via well-behaved updates. The formalization is as follows.

Definition 3.9 (Reachability subschemata and subviews). For $(N_1, N_2) \in \mathsf{Compat}\langle \Gamma_1;\Gamma_2\rangle$, define $\mathsf{Reachable}_\forall\langle \Gamma_1 : N_1 \parallel \Gamma_2 : N_2\rangle =$

$$\{M \in \mathsf{LDB}(\mathbf{D}) \mid ((N_1, N_2), (\gamma_1(M), \gamma_2(M)) \in \mathsf{IndUpd}_\forall\langle \Gamma_1\|\Gamma_2\rangle)\}.$$

Thus, $\mathsf{Reachable}_\forall\langle \Gamma_1 : N_1 \parallel \Gamma_2 : N_2\rangle$ is the set of all states of $\mathbf{D}$ which can be reached via independent updates on Γ_1 and Γ_2 from a state $M_0 \in \mathsf{LDB}(\mathbf{D})$ with $\gamma_1(M_0) = N_1$ and $\gamma_2(M_0) = N_2$. If $\{\Gamma_1, \Gamma_2\}$ forms a complementary pair, then this initial M_0 is determined completely by (N_1, N_2), but it is not necessary to enforce complementation in that which follows.

A limited view based upon Γ_1, which only involves the reachable states, is defined as follows.

(a) Define $\mathsf{Restr}_{\mathbf{D}}\langle \Gamma_1 : N_1 \parallel \Gamma_2 : N_2\rangle$ to be the subschema of $\mathbf{D}$ with $\mathsf{LDB}(\mathsf{Restr}_{\mathbf{D}}\langle \Gamma_1 : N_1 \parallel \Gamma_2 : N_2\rangle) = \mathsf{Reachable}_\forall\langle \Gamma_1 : N_1 \parallel \Gamma_2 : N_2\rangle$.

Thus, $\mathsf{Restr}_{\mathbf{D}}\langle \Gamma_1 : N_1 \parallel \Gamma_2 : N_2\rangle$ is the schema consisting of just those states reachable from (N_1, N_2). The corresponding sets of view states are defined as follows.

(b) For $i \in \{1, 2\}$, define $\mathsf{Restr}_{\mathbf{V}_i}\langle \Gamma_1 : N_1 \parallel \Gamma_2 : N_2\rangle = \{\gamma_i(M) \mid M \in \mathsf{Reachable}_\forall\langle \Gamma_1 : N_1 \parallel \Gamma_2 : N_2\rangle\}$.

The corresponding view morphism is then the appropriate restriction of γ_i.

(c) For $i \in \{1, 2\}$, define the function

$$\mathsf{Restr}_{\gamma_i}\langle \gamma_1 : N_1 \parallel \gamma_2 : N_2\rangle : \mathsf{LDB}(\mathsf{Restr}_{\mathbf{D}}\langle \Gamma_1 : N_1 \parallel \Gamma_2 : N_2\rangle) \to \mathsf{LDB}(\mathsf{Restr}_{\mathbf{V}_1}\langle \Gamma_1 : N_1 \parallel \Gamma_2 : N_2\rangle)$$

to be the restriction of γ_i to $\mathsf{LDB}(\mathsf{Restr}_{\mathbf{D}}\langle \Gamma_1 : N_1 \parallel \Gamma_2 : N_2\rangle)$.

Finally, the restricted view is obtained by assembling these pieces.

(d) For $i \in \{1, 2\}$, define

$$\mathsf{Restr}_{\Gamma_i}\langle \Gamma_1 : N_1 \parallel \Gamma_2 : N_2\rangle = (\mathsf{Restr}_{\mathbf{V}_i}\langle \Gamma_1 : N_1 \parallel \Gamma_2 : N_2\rangle, \mathsf{Restr}_{\gamma_i}\langle \Gamma_1 : N_1 \parallel \Gamma_2 : N_2\rangle)$$

to be the view of $\mathsf{Restr}_{\mathbf{D}}\langle \Gamma_1 : N_1 \parallel \Gamma_2 : N_2\rangle$ constructed from these.

That this view provides exactly that which is needed to support the extended and reversible set of independent updates for a pair of views is recaptured in the following.

Theorem 3.10 (The restricted view defined by a compatible pair). *The view* $\{\mathsf{Restr}_{\Gamma_1}\langle \Gamma_1 : N_1 \parallel \Gamma_2 : N_2\rangle, \mathsf{Restr}_{\Gamma_2}\langle \Gamma_1 : N_1 \parallel \Gamma_2 : N_2\rangle\}$ *has commuting congruences with*

$$\mathsf{Restr}_{\Gamma_1}\langle \Gamma_1 : N_1 \parallel \Gamma_2 : N_2\rangle \wedge \mathsf{Restr}_{\Gamma_2}\langle \Gamma_1 : N_1 \parallel \Gamma_2 : N_2\rangle = \mathsf{ZView}_{\mathsf{Restr}_{\mathbf{D}}\langle \Gamma_1 : N_1 \parallel \Gamma_2 : N_2\rangle}.$$

Proof. There is really nothing difficult to prove; the given properties are crafted right into the definition. In particular, the meet is the zero view because the interdependence conditions which place limitations on the allowable updates are enforced by including only those states which are already compatible. □

Examples 3.11 (Independent view updates in the reachability context). Consider first the views $\Pi^{\mathbf{E}_2}_{AB}$ and $\Pi^{\mathbf{E}_2}_{BC}$ associated with the schema $\mathbf{E}_2$, as introduced in Examples 3.4. Let $(N_1, N_2) \in \mathsf{Compat}\langle \Pi^{\mathbf{E}_2}_{AB}; \Pi^{\mathbf{E}_2}_{BC}\rangle$. The key information which is used to characterized the admissible updates is found in the sets $\pi_B(N_1)$ and $\pi_B(N_2)$. Specifically, $\mathsf{LDB}(\mathsf{Restr}_{\mathbf{E}_2}\langle \Pi^{\mathbf{E}_2}_{AB}\!:\!N_1 \parallel \Pi^{\mathbf{E}_2}_{BC}\!:\!N_2\rangle)$ $= \mathsf{Reachable}_{\forall}\langle \Pi^{\mathbf{E}_2}_{AB}\!:\!N_1 \parallel \Pi^{\mathbf{E}_2}_{BC}\!:\!N_2\rangle = \{(N'_1, N'_2) \in \mathsf{LDB}(\mathbf{E}_2^{AB}) \times \mathsf{LDB}(\mathbf{E}_2^{BC}) \mid \pi_B(N'_1) \subseteq \pi_B(N_1) \text{ and } \pi_B(N_2) \subseteq \pi_B(N'_2)\}$. Thus, updates to the schemata of these two views are constrained only in that the initial projection of the relation of $R_{AB}[AB]$ onto B may not increase, and the initial projection of the relation of $R_{BC}[BC]$ may not decrease. This is far more flexible than the constant-complement solution suggested in [15, Discussion 3.1]. In that solution, in order to maintain a meet situation, a copy of the projection of $\Pi^{\mathbf{E}_2}_{BC}$ onto B must be included in the view $\Pi^{\mathbf{E}_2}_{AB}$. That limits the allowable updates to those which keeps both $\pi_B(N_1)$ and $\pi_B(N_2)$ constant, a much smaller set.

Next, consider the views associated with $\mathbf{E}_1$. Here the classical constant-complement update strategy would allow no updates at all to either view. However, with the restricted views, the allowable updates are those which satisfy the conditions identified in Sec. 1. For a given $(N_1, N_2) \in \mathsf{Compat}\langle \Pi^{\mathbf{E}_1}_{AB}; \Pi^{\mathbf{E}_1}_{BC}\rangle$, using the definitions of $\equiv^{N}_{\langle X,Y\rangle}$ given in Sec. 1, $\mathsf{LDB}(\mathsf{Restr}_{\mathbf{E}_2}\langle \Pi^{\mathbf{E}_2}_{AB}\!:\!N_1 \parallel \Pi^{\mathbf{E}_2}_{BC}\!:\!N_2\rangle)$ $= \mathsf{Reachable}_{\forall}\langle \Pi^{\mathbf{E}_1}_{AB}\!:\!N_1 \parallel \Pi^{\mathbf{E}_2}_{BC}\!:\!N_2\rangle = \{(N'_1, N'_2) \in \mathsf{LDB}(\mathbf{E}_1^{AB}) \times \mathsf{LDB}(\mathbf{E}_2^{BC}) \mid \pi_B(N_1) = \pi_B(N_2) \text{ and } \equiv^{N'_1}_{\langle B,A\rangle} \subseteq \equiv^{N_1}_{\langle B,A\rangle} \text{ and } \equiv^{N_2}_{\langle B,C\rangle} \subseteq \equiv^{N'_2}_{\langle B,C\rangle}\}$. Parallel updates by the two views may reach any of these states.

The price paid for using this type of update strategy is that the constraints on which updates are allowed must be reset every time the pair of views is updated outside of this framework. That would happen, for example, when an update not supported in the restricted strategy were necessary, and so the two views would be combined, the update performed, and then a new initial compatible pair obtained. However, for many applications, this seems like a small price to pay in return for a substantially enlarged set of admissible independent updates.

4 Conclusions and Further Directions

A way to handle updates on two views, without any conflict, has been presented. The approach extends the classical constant-complement strategy in two ways. First and foremost, it is not restricted to meet complements (translatable strategies in the language of [6]). Rather, it takes advantage of the fact that simultaneous updates are limited in scope, and assumes that the updates to the companion view follow the associated protocol. Second, it does not depend upon complementation in any way, and so is readily extensible to any finite number of views.

Directions for additional investigation include the following:

Extension to finite sets of views: As noted in Sec. 1, a primary motivation for this work is the modelling in the context of many views. It is therefore of primary importance to develop the details of how this approach extends to more than two views.

Integration with applications: The ideas developed here should be of great use in extending the notion of database schema components, as described in [13] and [16], as well as their applications in update via cooperation [17] and objects for transaction [15]. The next task is to examine the details of such applications.

Effective methods for identifying the restricted state set: In the approach developed in this paper, the allowable updates are defined by a starting context (the reachability subschema). It is important to identify ways to characterize and compute effectively this context for classes of views which arise in practice.

References

1. M. Arenas, L. E. Bertossi, and J. Chomicki. Answer sets for consistent query answering in inconsistent databases. *Theory and Practice of Logic Programming*, 3(4-5):393–424, 2003.
2. O. Arieli, M. Denecker, and M. Bruynooghe. Distance semantics for database repair. *Ann. Math. Artif. Intell.*, 50(3-4):389–415, 2007.
3. O. Arieli, M. Denecker, B. V. Nuffelen, and M. Bruynooghe. Computational methods for database repair by signed formulae. *Ann. Math. Artif. Intell.*, 46(1-2):4–37, 2006.
4. R. W. Baldwin. Naming and grouping privileges to simplify security management in large databases. In *Proc. 1990 IEEE Symposium on Research in Security and Privacy*, pages 116–132. IEEE Computer Society Press, 1990.
5. F. Bancilhon and N. Spyratos. Independent components of databases. In *Proceedings of the Seventh International Conference on Very Large Data Bases*, pages 398–408, 1981.
6. F. Bancilhon and N. Spyratos. Update semantics of relational views. *ACM Trans. Database Systems*, 6:557–575, 1981.
7. F. Bentayeb and D. Laurent. Inversion de l'algèbre relationnelle et mises à jour. Technical Report 97-9, Université d'Orléans, LIFO, 1997.
8. F. Bentayeb and D. Laurent. View updates translations in relational databases. In *Proc. DEXA '98, Vienna, Sept. 24-28, 1998*, pages 322–331, 1998.
9. B. A. Davey and H. A. Priestly. *Introduction to Lattices and Order.* Cambridge University Press, second edition, 2002.
10. U. Dayal and P. A. Bernstein. On the correct translation of update operations on relational views. *ACM Trans. Database Systems*, 8(3):381–416, 1982.
11. S. J. Hegner. An order-based theory of updates for closed database views. *Ann. Math. Art. Intell.*, 40:63–125, 2004.
12. S. J. Hegner. Information-based distance measures and the canonical reflection of view updates. Technical Report 0805, Institut für Informatik, Christian-Albrechts-Universität zu Kiel, October 2008. An updated and corrected version, which will appear in *Ann. Math. Art. Intell.*, is available on the Web site of the author.

13. S. J. Hegner. A model of database components and their interconnection based upon communicating views. In H. Jakkola, Y. Kiyoki, and T. Tokuda, editors, *Information Modelling and Knowledge Systems XIX*, Frontiers in Artificial Intelligence and Applications, pages 79–100. IOS Press, 2008.
14. S. J. Hegner. Semantic bijectivity and the uniqueness of constant-complement updates in the relatiional context. In K.-D. Schewe and B. Thalheim, editors, *International Workshop on Semantics in Data and Knowledge Bases, SDKB 2008, Nantes, France, March 29, 2008, Proceedings*, volume 4925 of *Lecture Notes in Computer Science*, pages 172–191. Springer-Verlag, 2008.
15. S. J. Hegner. A model of independence and overlap for transactions on database schemata. In B. Catania, M. Ivanović, and B. Thalheim, editors, *Advances in Databases and Information Systems, 14th East European Conference, ADBIS 2010, Novi Sad, Serbia, September 20-24, 2010, Proceedings*, volume 6295 of *Lecture Notes in Computer Science*, pages 209–223. Springer-Verlag, 2010.
16. S. J. Hegner. A simple model of negotiation for cooperative updates on database schema components. In Y. Kiyoki, T. Tokuda, A. Heimbürger, H. Jaakkola, and N. Yoshida., editors, *Frontiers in Artificial Intelligence and Applications XX11*, 2011. in press.
17. S. J. Hegner and P. Schmidt. Update support for database views via cooperation. In Y. Ioannis, B. Novikov, and B. Rachev, editors, *Advances in Databases and Information Systems, 11th East European Conference, ADBIS 2007, Varna, Bulgaria, September 29 - October 3, 2007, Proceedings*, volume 4690 of *Lecture Notes in Computer Science*, pages 98–113. Springer-Verlag, 2007.
18. A. M. Keller. *Updating Relational Databases through Views*. PhD thesis, Stanford University, 1985.
19. R. Langerak. View updates in relational databases with an independent scheme. *ACM Trans. Database Systems*, 15(1):40–66, 1990.
20. O. Ore. Theory of equivalence relations. *Duke Math. J.*, 9:573–627, 1942.
21. S. L. Osborn and Y. Guo. Modeling users in role-based access control. In *ACM Workshop on Role-Based Access Control*, pages 31–37, 2000.
22. J. Rissanen. Independent components of relations. *ACM Trans. Database Systems*, 2(4):317–325, 1977.
23. K.-D. Schewe and B. Thalheim. Component-driven engineering of database applications. In *APCCM '06: Proceedings of the 3rd Asia-Pacific conference on Conceptual modelling*, pages 105–114, Darlinghurst, Australia, 2006. Australian Computer Society, Inc.
24. B. Thalheim. Component development and construction for database design. *Data Knowl. Eng.*, 54(1):77–95, 2005.